

Chapter I

Testimony: The Problem

This chapter defines testimony and sets out the framework we will use to analyze its epistemology and normative structure. Section 1 distinguishes testimony from other assertive speech acts. Section 2 presents an analysis of testimony derived from Grice's definition of non-natural meaning (Grice 1957). This Gricean analysis is not meant as a reduction of the concept of testimony to concepts that can be defined independently of testimony. Rather it gives a circular generalization concerning testimony that will help illuminate its epistemology and normative structure, even though it does not define testimony.

The Gricean analysis brings out the importance of testimony's epistemological status. Section 3 briefly states the question this raises: To what extent are we justified in believing what we are told, and is this justification (when present) evidential? This dissertation will be devoted to answering this question and showing how it is related to the question of testimony's normative status: What kind of responsibility does the teller have for her testimony?

1. What Is Testimony?

To begin, we must clear up our terminology, and distinguish our technical uses of "telling" and "testimony" from their colloquial uses. I will use "testimony" and "telling" as cognates, so that when Alice tells Sarah that it is raining she has presented Sarah with a piece of testimony. In ordinary speech, "tell" has many uses with many different

grammars, covering cases that are not even speech acts (“I could tell that she was pleased”) nor assertions (“I told her to get some rest”). “Testimony” ordinarily is used narrowly, suggesting a formal or legal setting. My use of “tell” and “testimony” will be intermediate. Unlike ordinary users of “tell,” I will refer only to selected assertive speech acts, but unlike ordinary users of “testimony,” I will refer to acts that do not take place in formal settings.

Of all the assertive speech acts, testimony in particular is an act that is presented as informative. The teller presents her testimony as something that itself could give the hearer a reason to believe what is told. Without this testimony, the hearer might lack any reason to believe; the testimony is not meant to depend on independent reasons that the teller has to believe what is said. As Karen Jones puts it (using “testifying” where we would use “telling”), “*testifying* that *p* contrasts with *arguing* that *p* insofar as it is the testifier herself who vouches for the truth of *p*; someone who argues that *p* lets the arguments vouch for themselves” (Jones 1999, p. 57). When a speaker gives an argument, the hearer’s reasons to believe should come from the argument’s content, and not from the very fact that the speaker is arguing for this particular conclusion. Thus the conclusions of this particular piece of philosophical writing will not be things that I am telling the reader: I mean them to be believed based on the arguments that are given, not because I vouch for them. On the other hand, if I write “This sentence is being written on a Wednesday morning,” I mean the reader to believe it because I say so; I have provided no reason for belief other than my vouching for it. I mean the reader to believe what I write by relying on my testimony.

In general, testimony is not meant to bring the hearer to realize something with resources that she already had available. If Sarah and Alice are both gazing out the window at the dripping branches, and Alice says, “It’s raining,” Alice’s utterance is not testimony. Alice does not mean her utterance to be informative; she is remarking on something that Sarah can see for herself. If, by contrast, Alice is standing at the window and Sarah is sitting where she cannot see outside, then Alice may be telling Sarah that it is raining. Alice may intend her say-so alone to convince Sarah that it is raining. This is the distinguishing characteristic of testimony: The teller means to be believed based on her say-so alone.

These considerations allow us to give a definition of testimony:

(Definition of Testimony) Testimony is the act of saying something to someone who may not already know or believe it, meaning her to take your say-so as a reason to believe it.

This definition of testimony will not quite be our final definition; at the very end of the dissertation (section VII.4) we will arrive at a definition of testimony that does not depend on whether the teller intends the speaker to take the testimony as a reason to believe. This will allow us to take into account cases of listless testimony, in which the teller does not particularly care whether the hearer believes her. Stating this more precise definition of testimony, however, will require an account of testimony’s normative structure (section VII.1), and this is not the time to give that account. The initial definition of testimony will be enough to derive the account of its normative structure (see Chapters IV and VII), which can then be used in the refined definition.

The current definition of testimony will be enough for the derivation of its normative structure because the current definition captures the central and paradigmatic instances of testimony. The instances that do not conform to this definition will be parasitic on those

that do; if people never intended their say-so to be taken as a reason to believe, there could be no cases of testimony that was not intended to be taken as a reason to believe. This definition will therefore be adequate until the very end of the dissertation. For instance, it will mark testimony out from other kinds of assertion, such as arguments, reminders, and jokes (if those be assertions).

It is, however, an exaggeration to say that every paradigm case of testimony is meant to be taken as a sufficient reason for belief. A teller may confirm a belief that the hearer already believes tentatively, intending that the hearer take the testimony as a reason to strengthen her belief. This contrasts with the case in which Alice says “It’s raining” to Sarah when both are looking out the window; that does not count as testimony because Sarah knows that it’s raining perfectly well both before and after Alice’s utterance, which is not meant to strengthen her belief. Conversely, the teller may realize that the hearer has so much evidence against what the teller says that the testimony alone will not provide a sufficient reason to change her mind. Nevertheless, she may intend her testimony to weaken the hearer’s contrary belief, because the hearer takes the teller’s say-so as providing some evidence for what is said (even though that evidence is outweighed by the evidence the hearer already had available). In each case, the teller means the hearer to take her say-so as *some* reason to believe, to strengthen an antecedent belief or weaken an antecedent contrary belief. To save space, I will omit these cases from the definition of testimony and subsequent discussion. When I speak of inducing belief, this should be taken to cover strengthening antecedent belief and weakening contrary belief.¹

The notion of telling just defined is pure in that it excludes statements for which any reason is given. When the speaker gives a reason to back up what she says, she has given

¹ Thanks to Joe Camp for pointing out the necessity of this proviso.

an argument; she means the hearer to believe her at least partly because of the reason that is given. For example, if Alice says to Sarah, “There’s a bittern in the garden; I heard it booming,” Sarah’s reason to believe that there is a bittern in the garden does not come only from the fact that Alice said there was. The fact that Alice heard booming provides Sarah with a reason to believe that there is a bittern in the garden that is distinct from Alice’s assertion *that there is a bittern in the garden*. So Alice’s assertion that there is a bittern in the garden is not testimony. Only an assertion presented with *no* extrinsic reason for belief counts as pure testimony.

Many assertions that are supported by reasons, though they do not count as pure testimony, nevertheless raise many of the same epistemological issues. When the speaker intends the hearer to come to a belief without taking the speaker’s say-so as a reason for belief, we may call the assertion *word-independent*; other assertions will be *word-dependent*. Testimony as defined is word-dependent, but many assertions for which reasons are given will also be word-dependent. Take the example in which Alice says to Sarah that there is a bittern in the garden, citing its booming. There are four possible cases, depending on whether Sarah depends on Alice’s word that she heard the booming and on whether Sarah depends on Alice’s word that the booming indicates a bittern. In three of these cases, Alice’s assertion is word-dependent, and Sarah’s reason to believe in the bittern is like her reason to believe testimony in the pure sense.

Suppose that Sarah is an expert on bitterns but is not present when the bittern booms. Alice later tells Sarah, “There’s a bittern in the garden; I heard it booming earlier,” perhaps accurately imitating the sound. Sarah can identify the booming as the booming of a bittern, but she depends on Alice’s word that Alice did hear the booming. In fact, “I

heard [this] booming earlier” is pure testimony, since Alice provides no evidence for it other than her say-so. A second case is when Sarah and Alice are both in the garden, but Sarah knows nothing of bitterns. Alice says, “There’s a bittern in the garden; hear that booming?” Here Sarah knows of the booming independently of Alice’s word, but she needs Alice’s word to know that the booming indicates a bittern. Alice is effectively telling Sarah, “Booming like that comes from bitterns.”² In a third case, Sarah neither hears the bittern first-hand nor knows that bitterns boom. Then, when Alice says “There’s a bittern in the garden; I heard it booming earlier,” Sarah is completely dependent on Alice’s word. She has much the same reason to believe what Alice says as if Alice had issued the pure testimony, “There’s a bittern in the garden.” (The reasons supplied by these utterances may not be exactly identical. If Alice is hard of hearing and cites the bittern’s booming, Sarah may have reason not to believe in the bittern.) Only in the fourth case, when Sarah hears the bittern and can identify bitterns by their booming, is Alice’s assertion word-independent. When Alice says, “There’s a bittern in the garden; hear that booming?” she is pointing out something that Sarah could have realized for herself with no prompting.

This dissertation will focus on pure testimony, for which no reasons are given. The idea is that such pure testimony presents in their most concentrated form the epistemological issues raised by word-dependent assertions. Pure testimony, for which no reasons are given, may be rarer than word-dependent assertions for which some

² It is not obviously possible to isolate a purely testimonial element in every analogous case, where the hearer depends on the speaker’s word to know that reasons given do support what the speaker says. For instance, Alice may not be committing herself to the claim that most booming comes from bitterns, or that most booming of a certain type comes from bitterns; it may be that most such booming comes from some other bird, but that Alice knows that there are no such other birds in the vicinity. The point is not to reduce such cases to cases of testimony, but to show that they involve similar epistemological issues of word-dependence.

reasons are given. There is no reason, however, to think that such word-dependent assertions will have a radically different epistemology from pure testimony. Most assertions for which reasons are provided will still be somewhat word-dependent, and analyzing testimony will shed light on all these cases of assertion.³ (We will briefly discuss reason-giving in section VII.2.)

In the remainder of this section, I will situate this definition of testimony with respect to other writers' conceptions of testimony and assertion in general. C.A.J. Coady, for instance, gives a narrower definition of testimony, which includes the requirement that the speaker "has the relevant competence, authority, or credentials to state truly [what she is telling]" (Coady 1992, p. 42). Our definition of testimony contains no requirement that the speaker be competent; a person may tell another something, intending to be believed on her say-so, whether or not she knows what she is talking about. Coady's definition needlessly complicates the epistemology of testimony; the hearer must determine whether the speaker has the relevant notion of competence in order to determine whether the utterance qualifies as testimony, and must then consider whether the testimony gives a reason for belief (even an authority might lie). As Elizabeth Fricker points out,

arguably it would be better to define the central concept in an epistemologically neutral way.... If the authority of the testifier were built into the definition of testimony, it might seem plausible that there is a general epistemic right to believe what is 'testified' to one; but this would simply pass the epistemic buck, from the hearer's point of view, to the question: is what I have just received a piece of 'testimony' or not? (Fricker 1995, pp. 396-7)

³ Since we discuss cases in which reasons are given, our analysis of testimony in isolation from the practice of giving reasons is not incompatible with a view such as Robert Brandom's, on which the practice of assertion requires a practice of giving and asking for reasons (Brandom 1994, pp. 170-1 and *passim*). Indeed, many of the assertions Brandom discusses will count as pure testimony on our definition. Brandom requires that the asserter be committed to giving a reason when challenged. If she is not actually challenged and does not actually give a reason, then her assertion may still be pure testimony.

Our definition of testimony, by contrast, makes testimony largely transparent: The hearer should in general be able to tell without difficulty when she is being told something, and what she is being told. As Fricker says, “a hearer competent in the speaker’s language is able to perceive that a speaker has, say, asserted that P when this is so” (Fricker 1987, p. 73); she is here using “assert” to refer to testimony in particular.

The transparency of testimony does have limits. Sometimes communication may break down: The hearer may misunderstand what the speaker is telling her, mistake a joke for serious testimony (or vice versa), or take a reminder for brute telling.⁴ When communication does not break down, however, it should be obvious to the hearer what she is being told, though it may not be obvious whether what she is being told is true. Thus, except in the extraordinary circumstances of breakdown of communication, the epistemological issues concerning testimony will focus on what the testimony gives the hearer a reason to believe, and not on what gives the hearer a reason to believe that she has heard testimony.

Our current definition of testimony, nevertheless, is not as transparent as it ought to be. The definition lacks transparency in that it refers to the teller’s intention to be believed. Even when communication does not break down, the hearer may not know whether the teller cares whether she believes her; the teller might simply want to be able to say “I told you so.” Our eventual refinement of the definition of testimony (section

⁴ These mistakes will be more likely the more indirect the speaker’s way of talking. If the speaker uses obscure idioms or private names, for instance, the hearer may mistake the content of what she is told. Irony can similarly contribute to mistaking non-testimony for testimony. If the speaker supposes that the hearer knows that she would never seriously assert something, she may say it as a joke; if, however, the hearer does not know that the speaker would never assert it, she may take it as sincere testimony. Both these cases are breakdowns in communication, which do not present serious problems for the idea that in normal cases testimony is transparent.

VII.4), which covers those cases in which the teller does not intend to be believed, will be fully transparent.⁵

The transparency of our current definition of testimony is adequate to defuse Paul Faulkner's attempt to broaden the notion of testimony to speech acts other than those that meet our definition: "if 'testimony' and 'presentation-as-true' are taken to be synonymous..., story telling and jokes are testimony" (Faulkner 2000, p. 586). Faulkner has in mind that story telling and jokes, like the informative acts we describe as "testimony," are communicative acts whose contents are truth-valued.⁶ Though this is a similarity among these speech acts, testimony as we have defined it forms a natural epistemological category because of its transparency. It should be clear to the hearer whether what has been said is presented as a joke or as a reason to believe; a speaker whose joke is taken seriously has failed to communicate, as much as one who has been heard to say "cat" when she meant "mat." The interesting epistemological question is not how the hearer knows that what is said is not a joke, but what reason she has to believe it once she knows it is not a joke. Thus, in exploring the epistemology of testimony, we will want to concentrate on those speech acts that are meant to provide a reason to believe, and we can set to one side other presentations-as-true.

Searle and Vanderveken present a list of assertive verbs that does not include "tell" (Searle and Vanderveken 1985, pp. 182-3ff.); our definition of testimony cuts across

⁵ See the discussion following the Definition of Testimony, above.

⁶ Faulkner's use of "testimony" to cover all presentations-as-true may be motivated by conformity to the usage of Tyler Burge (Burge 1993). Our discussion of Burge in Chapter II will not depend on Burge's account of presentation-as-true in general.

several of the verbs they list.⁷ The acts that Searle and Vanderveken call informing, predicting, reporting, and retrodicting (asserting a past proposition based on present evidence) may all count as testimony in our sense; they may all be meant to give the hearer a reason to believe based on the speaker's say-so. Our notion of testimony, however, delineates the natural category for dealing with epistemological questions. There may be epistemological differences between reporting and predicting; an informant may be better at saying what she has seen than at predicting the future. These differences, however, will only show up at a level of analysis that requires us to consider factors other than the speech act being performed, such as the topic of the testimony (see Chapter VI). On a less specific analysis, these different assertive acts will all have the same epistemological significance, so long as they are all acts of testimony (see Chapter IV).

Not all assertive acts, of course, count as testimony. For instance, reminders and arguments, which count as assertion, have an epistemological significance much different from that of testimony. Accordingly, our account of the epistemology and normative structure of testimony will not extend to these other assertive acts. This separates our account from that of Alston (2000), who attempts to give a unified account of the normative significance of assertion: The asserter is always responsible for the truth of her testimony (and must explicitly present its content, unless she is asserting it indirectly). Alston does not derive this normative structure from the epistemology of assertion, so the variation in the epistemological status of assertions is not a problem for his argument. Since we will derive the normative status of testimony from its epistemology (see

⁷ Searle and Vanderveken do list “testify” as a verb, but this applies to formal testimony under oath as opposed to the broader sense in which we use “testimony,” as cognate with “tell” (Searle and Vanderveken 1985, pp. 188-9).

Chapters IV and VII), we must restrict our attention to the epistemologically unified category of testimony, even if we think that in other assertive acts the speaker is also responsible for the truth of her assertion.

2. A Gricean Analysis of Telling

To illuminate the epistemology and normative status of telling, we will use an analysis that derives from Grice's account of non-natural meaning. Our use of the Gricean mechanism will be much less ambitious than Grice's original account (Grice 1957), which aspired to cover all informative/descriptive cases of meaning something by words, and which may have been meant to provide a reductive definition of meaning in terms of a certain kind of intention. Our Gricean analysis is meant to apply only to the particular act of telling, and is presented not as a reduction but as a generalization that holds true of telling. We will not attempt to reduce the concept of telling to other concepts.⁸

Grice's original account of meaning in informative/descriptive cases is:

[Meaning-NN] '*A* meant_{NN} something by *x*' is roughly equivalent to '*A* uttered *x* with the intention of inducing a belief by means of the recognition of this intention' (Grice 1957, p. 219).

As an account of all informative/descriptive cases of meaning something by words, this fails. There are many counterexamples in which a person obviously means that *p* without having the intention Grice describes; these eventually led Grice to make substantial modifications to his account (see Grice 1969). Rather than attempting to modify this definition to cover all cases of assertion, we will adapt it to be suitable as an analysis of

⁸ It is best not to see the definition of testimony in section 1 as a reduction to the concepts used in the definiens. If this definition were taken as reductive, we would be vulnerable to Dummett's criticism, discussed below, that belief that *p* cannot be defined independently of assertion that *p*.

telling in particular. This will avoid many of the counterexamples to Grice's definition; those counterexamples are cases in which the speaker means that p but does not tell anyone that p .

Unaltered, Grice's account will not even yield a true generalization concerning telling. We must make several changes, which I will explain sequentially. First, to clarify the changes, we will insert argument places for the hearer and for the content of the utterance, and suppress x , the uttered sentence, which will not concern us.

[Meaning*] A meant (to B) that p if and only if A intended to induce B to believe that p by means of the recognition of this very intention.

Of course this is still an account of meaning rather than telling; hence the unidiomatic construction "meant (to B)," since meanings, unlike tellings, are not addressed to specific people. To convert this into an account of telling, we will change "mean" to "tell" and "intend" to "attempt."⁹

(T-prelim) A told B that p if and only if A attempted to induce in B a belief that p by means of B 's recognition of that very attempt.

(T-prelim), however, is not satisfactory even as a generalization about the subclass of meanings that are tellings. It is vulnerable to a recondite counterexample introduced by P.F. Strawson, in which A attempts to induce a belief in B as described in the right side of (T-prelim), but A is not telling B anything. The counterexample is based on the possibility that A may intend the following: that B recognize that A intends to get her to

⁹ Though very little rides on it, the substitution of "attempt" for "intention" is not arbitrary. The idea is that "attempt" and "tell" are both *agentive* verbs, in Zeno Vendler's terms (Vendler 1957), while "mean" and "intend" are *stative* verbs. For further discussion of this sort of difference, see Mourelatos (1978, 1993) and Gill (1993). Again, very little rides on this; if we were to say that the teller intended to induce a belief rather than attempting to induce a belief, our subsequent arguments would be almost entirely unchanged.

believe something, but that *B* not recognize that *A* intends her to recognize this intention to get her to believe.¹⁰ The idea is that *A*

arranges convincing-looking ‘evidence’ that *p*, in a place where [*B*] is bound to see it. He does this, knowing that *B* is watching him [*A*] at work, but *knowing also that B does not know that A knows that B is watching him at work*. He realizes that *B* will not take the *arranged* ‘evidence’ as genuine or natural evidence that *p*, but realizes, and indeed intends, that *B* will take his arranging of it as grounds for thinking that he, *A*, intends to induce in *B* the belief that *p*. (Strawson 1964, pp. 446-7; I have changed Strawson’s variables; italics in original.)

The example is completed by the stipulation that (just as *A* intends) *B* will take *A*’s attempt to induce the belief that *p* as a reason to believe that *p*—she might, for instance, believe that *A* would not deceive her about the truth of *p*.¹¹ *A*’s act here satisfies (T-prelim): it is an attempt to get *B* to believe that *p* by means of *B*’s recognition of that attempt. Yet it does not seem to be an act of telling, or of any sort of communication; for instance, *A* does not intend that *B* recognize that any sort of communication has taken place.¹² (*A* intends *B* to think that *A* has attempted to play a benign trick on her.)

¹⁰ Joe Camp has suggested that this is not truly a counterexample, if (T-prelim) requires that *A* intend *B* to recognize that *A* intends to induce a belief in *B* by means of *B*’s recognition of the attempt to induce a belief by means of *B*’s recognition of that attempt, rather than by means of the recognition of the attempt to induce a belief, full stop. Adopting this reading of (T-prelim) would essentially yield (T), below. Grice himself did not hold to an analogous reading; he did take Strawson’s case as a counterexample (Grice 1969, p. 95).

¹¹ An example: Alice wishes to convince Sarah that the toilet tank is leaking, so that Sarah will call a plumber. She knows that Sarah is watching her, though Sarah does not know that she knows this. Alice scatters water on the floor near the toilet, reasoning as follows: “Sarah will think that I meant her to take this water as evidence that the toilet tank is leaking. She will thus recognize that I am attempting to get her to believe that the tank is leaking. She will then realize that I have no incentive to convince her that the tank is leaking (and to call the plumber) unless the tank truly is leaking (and the plumber is needed). So her recognition that I am attempting to induce the belief that the tank is leaking will lead her to believe that the tank is leaking.”

¹² Edward Hinchman (1998) argues that, in certain cases of conversational indirection, telling can take place without the speaker intending the hearer to recognize the speaker’s intention to communicate. Hinchman’s most convincing examples of this, for instance examples of insinuation, seem not to fit our definition of telling (Section 1). When I insinuate that *p*, I do not intend you to believe that *p* on my say-so; I intend you to draw the conclusion that *p* for yourself (perhaps I tell you that *q*, from which I intend you to conclude that *p* through a more or less tortuous process of reasoning). Thus these examples do not show that a speaker can tell someone something, by our definition, without intending the hearer to recognize that communication has taken place. (Hinchman does not see Strawson’s example as a case of telling.)

Gricean accounts often respond to this counterexample by building in an extra level of intentions or attempts: The teller must “*intend [B] to recognize his intention to get [B] to recognize his intention to get [B] to think that p* ” (Strawson 1964, p. 447). Even this extra level will not rule out all conceivable counterexamples; it may be possible to construct a counterexample in which the performer has the italicized intention just quoted, but does not intend the observer to recognize *that* intention. In theory a complete account could require an infinite number of levels of intention, although in practice after three or four levels it becomes impossible to construct a counterexample with any plausibility. Requiring an infinity of levels of intention and recognition would be unacceptable, since some of the higher-order intentions would likely be beyond our mental capacity. Alternative accounts have been proposed that preserve the reflexive structure without requiring an infinite hierarchy of reflexive intentions: see Sperber and Wilson on “mutual manifestness” (1986, pp. 38-46) and Bennett on “crossing” (1976, pp. 126-7). Such accounts avoid circularity in the Gricean definition.¹³ Since we do not seek a reductive Gricean definition, we have no reason to avoid circularity. This makes possible a more straightforward approach than Sperber and Wilson’s or Bennett’s.

This approach is to stipulate that the speaker intends the hearer to recognize that she has been told that p ; this will be built into the effect that the speaker is attempting to produce.

(T) *A* told *B* that p if and only if *A* attempted to induce in *B* a belief that p by means of *B*’s recognizing that *A* had told her that p .

¹³ We will not discuss the details of these accounts. They may be vulnerable to Dummett’s objection to the Gricean tradition (discussed below). For our purposes, their flaw is their complexity, in comparison to (T), below.

(T) is an analysis of the acts of telling picked out by the definition given in section 1: those acts that are meant to induce belief by means of the teller's say-so.¹⁴ Sometimes we will refer to the sort of attempt specified in (T) as a *T-attempt* (compare Grice's *M-intending*, Grice 1969, p. 105). (T) and the definition of testimony will in the end be replaced by an alternative analysis, which will be able to account for those cases in which someone tells another person something without particularly caring whether she is believed. This analysis will also flesh out the vague phrase "by means of," specifying exactly how *A* intends *B*'s recognition to bring about a belief. This refinement of (T), however, will have to wait until we have established the normative structure of testimony (see section VII.4). In the meantime, (T) will help us investigate testimony's epistemology and normative structure, and (like our definition of testimony) it will capture the central and paradigmatic cases of telling, in which the speaker does intend to be believed.

Note that (T) is explicitly circular. It would be no use in an effort to define testimony in terms of other, antecedently understood concepts. Nor is it intended to contribute to such a reductive enterprise. Rather it is a generalization that holds roughly of the acts of testimony picked out by our definition (Section 1), and that will give us a fruitful perspective on the epistemology and normative structure of testimony. As such, its circularity is no more vicious than the circularity of "x is human if and only if x has two human parents"; that does not provide a reductive definition of the concept "human," but it does tell us something useful about that concept. (T) also requires that both *A* and *B* have the concept of "telling," since *A* intends *B* to recognize that the *A* has told her that *p*.

¹⁴ Including acts that are meant to strengthen an antecedent belief or to weaken a contrary belief; see section 1.

Plausibly, anyone who is capable of giving testimony or of gaining justification through testimony has the concept of “telling,” so this intentional circularity is also benign.

Viewing (T) as a useful generalization rather than as a reductive definition avoids Michael Dummett’s charge against the Gricean tradition, that it explains assertions in terms of the thoughts they assert and must “take[] as already given the conception of the thought expressed by a sentence” (Dummett 1993, p. 172). This is fatal, according to Dummett, because we cannot understand what thoughts are without understanding how sentences express them. If we wanted to make something like (T-prelim) the basis of a reductive definition of telling, we would indeed need “belief that p ” to be definable independently of “telling that p ,” on pain of vicious circularity in the definition. This would at least arguably require what Dummett objects to, an understanding of the thought that p that was independent of any understanding the use of a sentence to express p . Since (T) is an admittedly circular generalization, however, we can use it even if we cannot understand its right side independently of an understanding of its left side. The use of (T) does not require that we understand what it is to believe that p independently of an understanding of what it is to tell someone that p .

The circularity of (T) helps this analysis of telling account for the desideratum of transparency of telling, mooted in Section 1: that, unless communication has broken down, the hearer should know that she has been told something and what she has been told. If telling is as described in (T), then part of what the teller is attempting is to make clear to the hearer that the teller has told her that p . If the hearer fails to recognize this, then the teller’s attempt must fail—even if the hearer comes to believe that p , she will not come to believe it by means of her recognition that the speaker has told her that p . Thus

the hearer's failure to recognize that she has been told that p is a communication breakdown that frustrates what the teller is attempting to accomplish.

Telling someone that p according to (T), then, counts on the fact that the hearer will generally be able to recognize that she has been told that p . This will leave the more difficult and interesting epistemological problem that we state in detail in the next section: Why should anyone believe anything that she is told?

3. Statement of the Epistemological Problem

Much of this dissertation investigates the epistemology of testimony. The question here is, when do people have reason to believe what they are told, and if they ever do, what kind of reason do they have? The answer to this epistemological question will allow us to derive an account of testimony's normative structure (Chapters IV and VII); indeed, as we will see, we will have to account for the norms on testimony in order to resolve some of the objections to our epistemological view (section III.3). The epistemological question, however, is important for more reasons than its effects on testimony's normative structure. It goes to the heart of what testimony is.

As we have defined testimony (Section 1) and analyzed it in (T) (Section 2), its epistemic effects are all-important. According to the definition, testimony is meant to serve as a reason to believe, and according to (T), testimony is a T-attempt to induce a belief. The question of whether testimony actually does give a reason for belief is critical for any understanding of how these attempts to induce belief might succeed. If no one

ever came to a belief because of testimony, testimony would be futile.¹⁵ Contrast betting. Bets are not epistemically inert; for instance, if Alice bets someone that it will rain, Sarah may conclude that Alice believes that it will rain, and this may give Sarah a reason to believe that it will rain. The practice of betting, however, does not depend on this epistemological status. It would make sense to bet even if no one ever did draw a theoretical conclusion from an act of betting; money could still change hands. The importance of the epistemic effects of testimony, by contrast, are built into its definition.

Our analysis of testimony's epistemological status will not address any questions about the psychological processes by which testimony actually produces or fails to produce beliefs. Though we will have to address some questions about the psychology of the teller (Chapter V), we are not concerned with the psychology or phenomenology of the hearer of testimony. Suppose Alice tells Sarah that it is raining; we are not concerned with whether Sarah performs mental calculations concerning Alice's reliability, or whether she simply hears Alice and believes her. Most beliefs are not based on reasoning that the believer explicitly mentally rehearses. A belief may be justified by reasons that the believer has not thought of, or even that she cannot produce on demand. This does not mean that any belief is justified so long as there is some argument that would justify it if thought of. A believer may blindly guess something that she would be justified in believing if she were sensitive to the evidence, but the blind guess will not be justified. Even if the believer cannot produce an argument for her belief, the belief will not be

¹⁵ This applies even to those cases, not covered by our definition, in which the teller is indifferent to whether she is believed; they are parasitic on cases in which testimony is meant to be believed. See Section 1.

justified unless the method that produced it is somehow sensitive to the considerations that justify it.¹⁶ Perhaps the believer should be seen as reasoning subconsciously.

What concerns us is akin to the logical positivists' "rational reconstruction," which makes sense of reasoning procedures without necessarily mirroring what goes on in the reasoner's mind. As Carl Hempel describes the rational reconstruction of scientific inquiry, it is "the construction of a consistent and comprehensive theoretical model of scientific inquiry, which is then to serve as a system of reference, or a standard, in the examination of any particular scientific research" (Hempel 1945, p. 44). The arguments modeled in the rational reconstruction are an idealized version of the conscious or unconscious thought processes that take place in the formation of beliefs; they give the justifications that might be available to the believer. So a rational belief need not be attained according to the arguments given in the rational reconstruction, so long as a rational reconstruction is available of how the belief actually was attained.¹⁷

In the case of testimony, we will be exhibiting justifications for the hearer's beliefs that advert to, for instance, the teller's past testimony. We need not claim that in these cases the hearer mentally rehearses what she knows about the teller's past testimony before deciding whether to believe her this time. Rather we need only claim that the hearer is sensitive to the teller's past testimony in that, if the teller had told more

¹⁶ This goes beyond the conception of justification set forth in the Appendix. The Appendix brackets the problem of distinguishing beliefs that are justifiable from beliefs that are actually justified, arrived at in the proper way. The considerations set forth in the text might be helpful in an attempt to deal with the problem of distinguishing justifiable beliefs from justified beliefs.

¹⁷ Hempel himself may have thought that a procedure that did not conform to the rational reconstruction was thereby erroneous (even if its results turned out to be correct or reliable). After the quoted passage, he compares rational reconstruction to determining the rules of a game from the moves that players make, allowing for the possibility that the players occasionally violate the rules of the game. Such a view of rational reconstruction, on which it determines rational procedures (which should be consciously followed) from actual procedures of thought that may be irrational, would discard too much of our thought as irrational, since we rarely do explicitly follow those procedures.

falsehoods in the past, the hearer would be less likely to believe her this time. Perhaps the teller's past reliability produces a feeling of security in the hearer as she accepts the testimony, whereas a teller of past falsehoods would produce a feeling of doubt, even if the hearer could not name the source of that doubt. What is important is that the teller's past reliability has some effect on the hearer's thought processes.

Our investigation, then, will focus on when testimony gives the hearer a justification for believing what she is told, and what kind of justification it can give. We will argue that particular pieces of testimony provide justification by providing evidence for what is told. (Chapter II explains why this claim must be restricted to particular pieces of testimony.) This is not the implausible claim that we are justified in accepting testimony only when we have explicitly weighed up the evidence for and against it; weighing up evidence would be part of a rational reconstruction of the hearer's thought processes.

Opposed to this view will be the Assurance View, held by Angus Ross (1986) and Richard Moran (1999). The Assurance View holds that testimony provides a non-evidential justification for belief, rooted in the teller's assumption of responsibility for her testimony (see chapter III). The burden of proof will be on the Assurance View to show that testimonial justification is non-evidential. Except for clear *a priori* justifications such as mathematical proofs, epistemically justified beliefs should generally be those that are (on a rational reconstruction) supported by the evidence. Anyone who wishes to posit a non-evidential justification will have to explain why it should be considered to be justification and why it should be considered not to be evidential. Otherwise we run the risk of positing mysterious non-evidential *a posteriori* justifications

for belief, without explaining what a non-evidential *a posteriori* justification for belief might be.

To say that the burden of proof is on the non-evidential view of testimony, however, is not to say that the burden cannot be met. Testimony's ubiquity and its status as an intentional act will create problems for the evidential view of testimony. These have to be addressed if we are to be able to see testimony as giving evidence. In the next chapter we will distinguish two senses in which we might say that we are non-evidentially justified in relying on testimony, and discuss the sense in which it is true that we have this non-evidential justification. In Chapter III we will discuss an argument that we have non-evidential justification in another sense; Chapters IV through VII then refute this argument, demonstrating the sense in which testimony does provide evidence.